

AUTONOMY FOR CONSTELLATIONS

Walt Truszkowski
Senior Technologist
Code 588
NASA - Goddard Space Flight Center
Greenbelt, MD 20771
Walt.Truszkowski@gsfc.nasa.gov

David Zoch, Dan Smith
Lockheed Martin Space Operations
7375 Executive Place, Suite 101
Seabrook, MD 20706
{david.zoch, dan.s.smith}@lmco.com

1. Introduction

A new class of planned space missions is placing challenging demands for improved autonomy concepts and techniques. Among the motivations for these challenges are resource constraints like onboard processor speeds, memory, s/c power, etc. Even though onboard computing power will surely increase in the coming years, the resource constraints associated with space-based processes will continue to be a major factor that needs to be addressed when dealing with, for example, agent-based spacecraft autonomy.

To realize "economical intelligence", i.e., constrained computational intelligence that can reside within a process under severe resource constraints (time, power, space, etc.), is a major goal for such space systems as the Nanosat constellations.

To begin addressing the new challenges, we are developing approaches to constellation autonomy with constraints in mind. Within the Agent Concepts Testbed (ACT) at the Goddard Space Flight Center (GSFC) we are currently developing a Nanosat-related prototype for the first of the two-step program detailed below.

The two-step program for constellation autonomy studies is as follows:

- (1) Develop a community of surrogate ground-based agents representing the satellites in the constellation. This will enable us to establish, in a prototype environment, the centralized and distributed agent behaviors that we postulate will be eventually used in space.
- (2) Migrate the surrogate community to simulated space-based satellites.

This paper will focus on Step 1 in the context of our ACT activities and present some ideas relating to Step 2. First we present a brief discussion of constellation mission models and their associated challenges. These challenges motivate the agent-based technology work we are pursuing with the goal of achieving constellation autonomy.

2. Constellation Mission Models

A wide variety of missions can be implemented best with constellations of satellites working together to meet a single objective. Reasons cited for using constellations include lower mission cost, the need for coordinated science, special coverage or survey requirements, and the need for quick-reaction tactical placement of multiple satellites.

The cost of producing the satellites and getting them to orbit may actually be lower than the traditional "one of a kind" satellites with a dedicated launch. With a traditional satellite, system reliability

requirements force a high level of component protections and redundancy, which leads to higher overall weight and launch costs. A dedicated launch is often required to for these missions. With a constellation, system reliability can be met by having spare satellites. The use of per-satellite redundancy can be significantly reduced. In some cases, it may be practical to use lower-rated components at a much lower cost combined with an on-orbit-sparing plan. Additional savings are obtained through the use of assembly-line production techniques and coordinated test plans so that the satellites can basically be mass-produced. With a reduced size and weight, new options are made available for satellite launches. Lower cost launch vehicles are available, multiple satellites of the constellation can be launched at once, or piggyback launch slots can be obtained, where launch costs are shared with another user.

A constellation of as few as two satellites can be used to perform coordinated science functions. Storms and other phenomena observed from multiple angles can be used to generate 3-D views. Satellites with a wide spatial separation can be used for parallax studies of distant objects. A cluster of satellites flying in formation and working together can form a virtual lens hundreds of miles across.

The most widely used application of satellite constellations is for reasons of extended area coverage. Low earth orbiting constellations such as Globalstar use dozens of satellites to provide continuous global or near-global coverage. The GPS system uses a constellation to provide global coverage and spatial diversity of the multiple satellites in view. Earth mapping missions can use multiple satellites to shorten the time between successive observations of the same area. NASA is planning missions of up to 100 satellites for magnetospheric research with orbits nearly as distant as the moon. A single-satellite approach would require many years of data collection to match what the constellation can survey in a short amount of time.

Military applications for constellations include earth observation, weather, and equipment resource monitoring. In the future, it may be possible to launch very small satellites with a very specific purpose and very short mission duration. The satellites could be produced by the hundreds and launched as needed. The “constellation”, at any point in time, would include those satellites which were recently launched and which are currently performing their intended function.

Table 1 summarizes applications of satellite constellations and some of the critical distinctions between their mission models.

Table 1. Future constellation missions vary widely in their mission requirements.

TYPE	APPLICATION	TYPICAL DESIGN/ MANUFACTURE	DATA ACQUISITION	OPERATIONS
Simple (number of satellites)	University sponsored, corporate R&D	Very low cost, earth-rated components	Not a major issue, low rate, may operate at amateur radio frequencies	Extremely low cost, university level
Cluster [Cluster II (4), Magnetospheric Multiscale (5)]	Coordinated science, virtual telescopes, stereo imaging	Complex, satellite crosslinks, extensive testing required, high redundancy within satellites	Not a major issue. Typically high rate due to science mission, but number of satellites is limited or downlink access can be controlled	Similar to a large single satellite, multiple satellites performing a single coordinated function, additional effort for mission planning
Coverage Constellation [Globalstar (48), Orbcmm (36), TIROS (5), NASA NanoSat (100)]	Commercial phone/paging/internet systems, earth observation, (multi-point data collection, broad survey or coverage)	Satellites operate independently, designed for mass production, limited redundancy	May be large number of satellites utilizing many ground sites concurrently. Dedicated antenna sites may be needed due to high duty cycle	May involve hundreds or thousands of passes per day. Ideal for automation, as there are many nearly identical satellites working
Military/Tactical [XSS-10, ESCORT, Orbital Express]	inspection, imagery	New concepts are for very small low cost mass-produced with no redundancy and minimal mission duration	Only a few satellites activated at a time, may use portable data acquisition sites. May have a video downlink plus minimal status info	Primarily an orbit/maneuver and data acquisition activity. Data is for immediate use only. No long-term trending, etc.

Efficient mission operations are critical to the success of satellite constellation missions. It is unlikely that either commercial constellation missions or NASA missions will be launched unless it is known that the operations costs for a constellation of “n” satellites is significantly less than “n” times the costs for one satellite. NASA has set a goal of operating a constellation for the same cost as a single large satellite.

Two key challenges for efficient operations are the number of satellites in the constellation and the total number of satellite contacts per day. A data base administrator may have 100 different databases to maintain. The real-time operations personnel may, at some level, monitor over 1000 passes per day. Each ground operations role, from data base administrator to off-line analyst, must take into account the multiplier placed on each activity. By analyzing the complete workflow, one can select the most critical areas for automation or new operations concepts development (See Table 2). These concepts have been developed with future constellations of 50 to 100 satellites in mind. However, the concepts can be scaled down to apply to much smaller constellations or even to single satellites. Traditional ground support systems designed for single satellite support generally do not efficiently scale up to handle large constellations.

Table 2. Efficient constellation operations require coordination across multiple disciplines.

SELECTED AREA OF CONCERN	ISSUE	RESOLUTION	RESULT
Database	Each satellite may have slight differences	Tools to support constellation-wide changes as well as satellite-specific changes	Nearly as efficient as single satellite system
Mission Planning/Scheduling	Large number of satellites, need for coordinated operations, balancing the load for operations	Requires efficient, flexible, constraint-based automated conflict resolution	Should require same staff as large single satellite mission
Pre-pass and post-pass reconfigurations	Even a 1 minute effort becomes large if there are 1,000+ passes per day	Automate per a schedule or data-driven rules, avoid need for human intervention	Just as efficient as single satellite system
Routine satellite monitoring	Typically very labor intensive, effort multiplied for constellations	Expert systems, better data presentation, data/schedule-driven operations	One operator should be able to monitor many satellites
off-line analysis	Many satellites to trend	Automated analysis, use of other satellites in constellation as “control” satellites for investigating anomalies	Still labor intensive, but highly automated. Use the high number of satellites to your advantage in anomaly analysis
Flight Dynamics	Labor and processing intensive, may need to manage a formation	Ideal place for onboard autonomy	Reduced impact over traditional single satellite approach
Data Capture	Concurrent passes, large data volume	Automated data acquisition system, expanded data mngt tools, onboard data reduction	Autonomously run system with minimal intervention

With a constellation of 60 satellites, functions, which would take a minute for a single satellite, would take an hour. A 30-minute daily scheduling run for a traditional single-satellite system may seem short, but if multiplied by 60 would be unworkable. Similarly, a 1-minute manual system configuration for a new satellite pass may seem short, but a large constellation could easily lead to 1500 passes per day with passes starting or stopping every 30 seconds. If the concept of “effort multipliers” is not adequately addressed through operations concept development and ground segment design, then the operational staff required to support the mission will also multiply, possibly to the point of being prohibitively expensive.

Simple automation of existing functions is not sufficient to address the efficiency concerns for large constellations. New operations approaches and new system designs are required to meet the many challenges. Mission requirements concerning mission-level reliability, data acquisition volume and timeliness, acceptable down-time duration's, and levels of onboard autonomy must all be carefully developed to balance mission functional, developmental cost, and operations cost objectives.

Satellite constellation missions now being considered cover a wide range of characteristics and it may appear that a common solution of ground operations is not feasible. Satellite systems will vary by orders of magnitude in their data rate and total data volume, orbits may vary from low LEO or to very elliptical with multi-day orbit periods, air-to-ground protocols will vary widely, and the satellites themselves may be low-cost with low autonomy or may be sophisticated with a high level of onboard self-management.

The interested reader may refer to [1,2,3,4,5,6] for material supporting this section.

The solutions being investigated at NASA/GSFC in the ACT provide a flexible architecture for addressing automation across a domain that includes the control center, data acquisition sites, and onboard the satellite. It is anticipated that this testbed will support the tools and techniques to meet the constellation challenges identified in Table 2.

3. Ground-based Constellation Autonomy

The following provides an overview of the prototype being developed in the ACT [7,8,9]. Figure 1 provides a high-level graphic representation of the prototype. The prototype was implemented to demonstrate the agents' plan/execute/monitor cycle of execution. Agents execute a plan that consists of plan steps. Once plans have been generated, the agent verifies that preconditions are met for the steps in the plan, executes the steps, monitors the results, and updates state information based on the results of each step.

The prototype consists of a number of agents connected to a simulation environment for the ground stations and satellites as described below.

3.1 Spacecraft

There are currently four simulated spacecraft in orbit collecting magnetosphere data. These constitute the demonstration's constellation. The simulation environment currently propagates the orbits based on ideal conditions. Faults can be inserted into the telemetry stream to simulate an anomaly.

3.2 Ground Stations

The ground station simulation (GSS) receives commands that contain configuration information, position the antenna, etc. The simulation responds to these commands and generates a data flow of simulated telemetry from the (simulated) s/c once it has detected Acquisition of Signal (AOS). Our prototype data acquisition architecture consists of two ground stations at opposite sides of the earth with one antenna each.

3.3 Agents

The following are the agents in the current prototype:

- Mission Manager Agent (MMA) – Coordinates the agent community in the Constellation Control Center (CCC).
- Contact Manager Agent (CMA) – Communicates with the spacecraft, sends and receives data, commands, and telemetry.

- Spacecraft Proxy Agents – (SPA) There is a proxy agent for each spacecraft. The agents keep track of spacecraft status, health and safety, etc. The agents will flag the Mission Manager Agent when an emergency arises that may need possible replanning.

The prototype also contains several visualization tools, such as the community visualization tool, and a tool for visualizing internal agent messaging. Acquisition schedules are created by an external tool / agent called the Planner / Scheduler Agent (PSA).

3.4 Scenario

The main steps in the demonstration scenario are as follows:

1. Agents register with the Mission Manager Agent at system startup.
2. MMA determines it is time to make a new acquisition schedule and notifies the Planner/Scheduler Agent
3. Planner/Scheduler Agent communicates with the S/C Proxy Agents to get view data, and the Contact Manager Agent to get resource data (e.g., availability). It then creates a contact schedule for all spacecraft.
4. The schedule reaches the Contact Manager Agent. The Contact Manager Agent contacts the spacecraft at the appropriate time, and notifies S/C proxy agents of where to get the telemetry. Telemetry is downloaded.
5. The S/C Proxy Agent processes the telemetry data, updating the spacecraft's status, and evaluates any anomalies.
6. The Contact Manager Agent ends the contact when scheduled (or the ground station simulation sends a LOS if this occurs first).

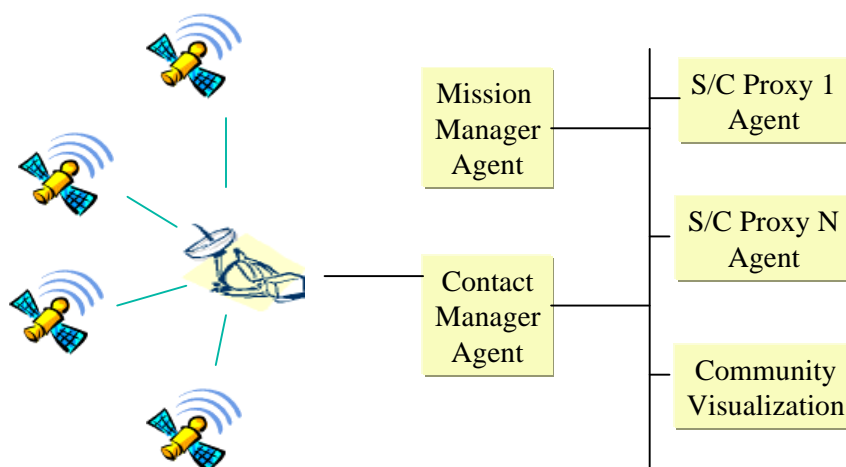


Figure 1. The ACT testbed consists of a community of cooperating Agents each of which is component-based

The group of surrogate spacecraft agents, as a major component of the ground-based community, maintains an awareness of the actual physical constellation. The surrogates act in behalf of their respective spacecraft in status monitoring, fault detection and correction, distributed planning and scheduling, and spacecraft cooperative behaviors (as needed).

A next step in the evolution of the ACT demonstration will be to have communities of agents each associated with a particular spacecraft in the constellation. Each of these communities will have specialist subsystem agents who will monitor the various subsystems of the spacecraft and cooperate with one another in the handling of anomalous situations. An overall coordinator, or spacecraft agent, will lead the community and represent the spacecraft to ground controllers. It will also represent the spacecraft to other spacecraft agents in the constellation community for activities such as distributed planning and scheduling, and other forms of collaboration. This approach is in keeping with Goddard's philosophy of dealing with a community of agents rather than a single monolithic agent.

In the context of spacecraft constellations, the ground-based group of surrogate agents serves in two differing capacities: (1) The prototype establishes the fact that surrogate agents can indeed support the concept of constellation autonomy in a meaningful way, and (2) having a ground-based community of surrogate agents allows developers and users (controllers) to gain confidence and trust in the approach.

The ground-based community provides invaluable experience in developing and evaluating agent-community support for constellation autonomy. This experience will more readily enable the migration of agent-community capabilities to the actual spacecraft in the constellation when the constellation can, in fact, support such a move.

4. Space-based Constellation Autonomy

The next step in the overall plan to realize space-based autonomy is to migrate the spacecraft surrogate community of agents to the actual spacecraft. This is a non-trivial step. A major step in the direction of actual onboard spacecraft autonomy is to have the agent community demonstrate its correctness in actual ground-based spacecraft control centers. This is further discussed in Section 5.

There are several issues that need to be addressed before this becomes a reality:

- Adaptation to resource constraints. As an example, a spacecraft subsystem agent must be able to exist and operate within the microprocessor associated with the subsystem. This is where the concept, which we call "economical intelligence", comes into play. Reasoning code and knowledge and information structures and management need to be "optimized" in order to function properly in the resource-constrained environment of a spacecraft subsystem microprocessor.
- Integration with existing subsystem autonomy. Currently, most spacecraft subsystems have a degree of autonomy already built into their operations. This is usually realized through the use of expert system or state-based technologies. A subsystem agent should be able to take advantage of the existing capability and build upon it. In this case the already existing capability would become an external resource to the subsystem agent that can be used to realize a higher level of autonomy for the subsystem. The agent needs to know about the external resource and how to use it, i.e., factor its information into its reasoning process.

- Real-time activity. Most situations on a spacecraft require real-time attention. If the situation is not readily handled by built-in subsystem autonomy the associated subsystem agent will need to respond in real-time. This will require the agent to have a working reflexive behavior.

Addressing these issues will be the foremost tasks in the next stage of our work in developing agent-based spacecraft autonomy.

5. Grand View

Figure 2 illustrates what we consider to be a grand view (not the only one) of what might happen in the distributed agent-based systems world. It paints a rich tapestry in which we can see many threads of agent-based activity both ground-based and space-based. The major theme of the figure is that of agent migration from one level to another. This is the theme of our approach to realizing spacecraft constellation autonomy. The figure shows three distinct levels: agent development, ground-based autonomy, and space-based autonomy.

Figure 2 also shows three levels of activity: agent development, ground-based autonomy, and space-based autonomy. The agent development is the initial stage during which agents are designed, implemented and verified. Prototype agent communities are also investigated at this level. The current ACT demonstration system exists at this development level of the agent migration hierarchy. Once confidence and trust in agent performance is achieved at this first level, the agents and agent communities can migrate to the ground-based autonomy level.

At the ground-based autonomy level we see agents supporting automated control centers as well as intelligent information management in information archives. Autonomous decision making allows to agents to migrate from node to node at this level of the hierarchy. At this level we also see the concept of spawning (or cloning). This capability allows for parallel activities (of the same kind) and brings about a type of fault tolerance in the case that an agent "goes down" for some reason. The concept of persistence is also pictured. This capability allows the agent to perform a task over an extended period of time.

The top level is space-based autonomy. Here we picture agents migrating to a spacecraft that is part of a constellation. Each agent is either associated with specific subsystems or performs the role as spacecraft agent having control or access to the subsystem agents. The spacecraft agent is the one that will maintain communication with the ground-based system reporting on overall spacecraft status or asking for help if a situation arises which the onboard community cannot handle. This agent will also be responsible for coordinating activities with the other spacecraft agents in the overall constellation.

The ultimate goal of the current ACT activities is to develop and verify agents and agent communities concepts to the point that they can migrate to actual ground-based operations and, when fully verified in this actual operational context, migrate successfully to provide onboard autonomy.

Figure 2 as well depicts the various migration paths that could be taken by agents and communities of agent's enroute to a spacecraft.

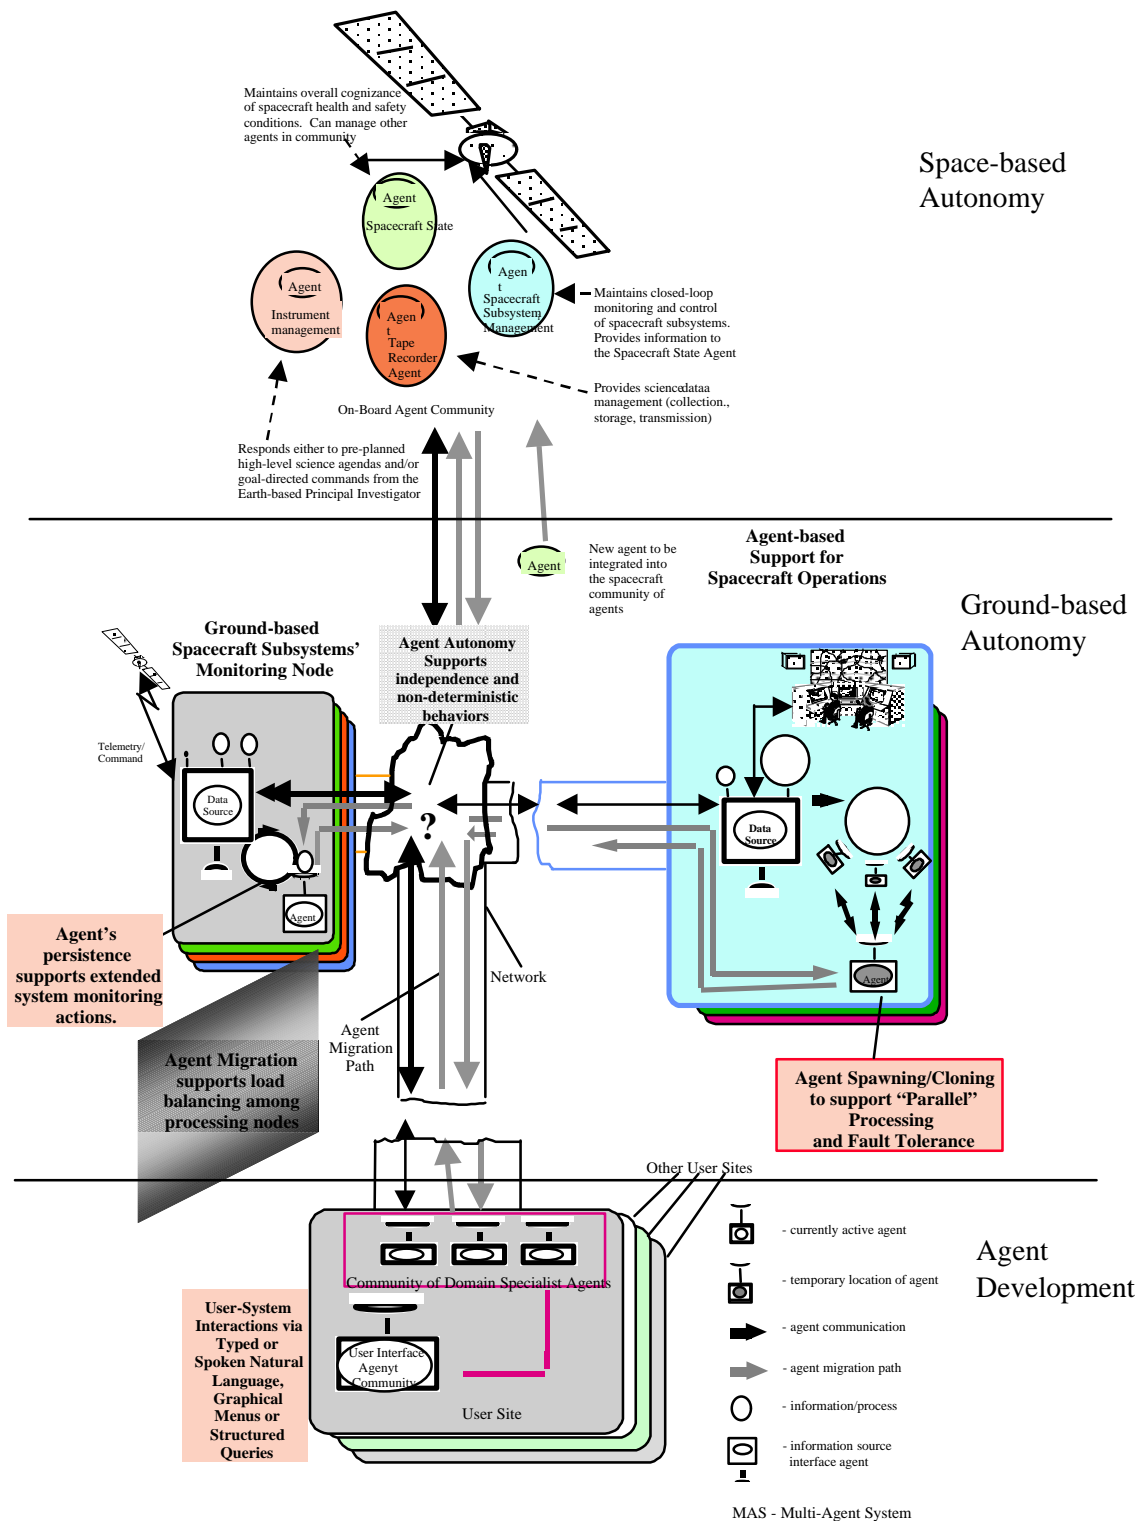


Figure 2. Agent Development and Deployment

6. Conclusion

This paper has addressed the issue of spacecraft constellation autonomy from the perspective of the agent technology research and development efforts going on in the context of the ACT. The idea of realizing constellation autonomy first through ground-based communities of spacecraft surrogate agents and then migrating the agent community technology to the actual spacecraft seems to be a reasonable approach. The progressive autonomy that could be realized through this approach will enable ground-based controllers to upload only those agents in the community that have been thoroughly verified and in which there is the needed trust.

7. Acknowledgements

The authors wish to thank the Goddard Agents Group for their work on the ACT demonstration system and the NASA SOMO Program for providing the funding.

References:

1. Daryl G. Boden and Wiley J. Larson, Cost-Effective Space Mission Operations, McGraw-Hill , 1996.
2. Stephen J. Talabac, Spacecraft Constellations : The Technological Challenges in the New Millenium, NASA/Goddard Space Flight Center, AETD Information Systems Center, September 1999.
3. Dan Smith, Efficient Mission Control for the 48-Satellite Globalstar Constellation, Third International Symposium on Space Mission Operations and Ground Data Systems.
4. James R. Wertz, et. al., Autonomous Orbit Control: Initial Flight Results from UoSAT-12, American Astronautical Society, AAS 00-71 Twenty-third Annual AAS Guidance and Control Conference.
5. Gary Meyers, Constellations: The new problems they present ... NASA/Goddard Space Flight Center, Presentation, 5/24/1999.
6. Rhoda Shaller Hornstein, John Willoughby, et. al., Cost Efficient Operations: Challenges from NASA Administrator and Lessons Learned from "Hunting Sacred Cows" SpaceOps '96.
7. Agent Architecture, 1999, Internal Goddard Report.
8. ACT Requirements Document, 1999, Internal Goddard Report.
9. ACT System Operations, 1999, Internal Goddard Report.